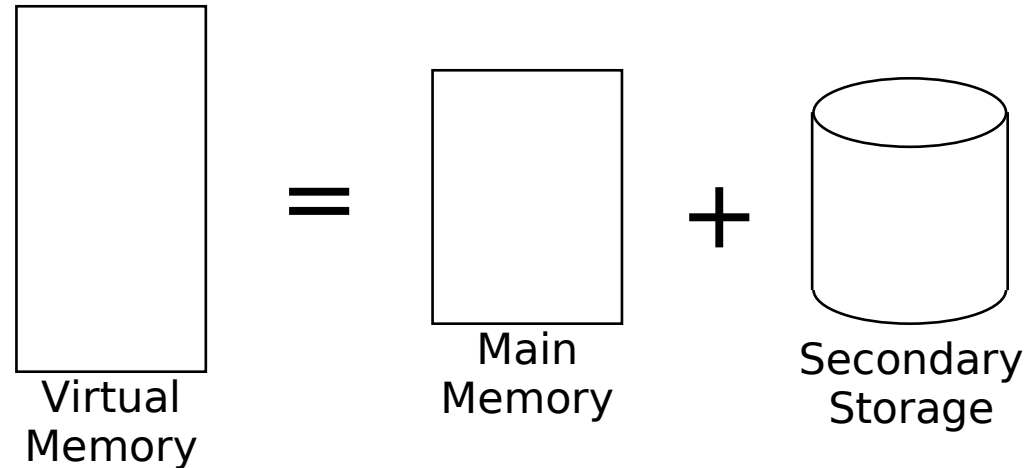


Virtual Memory



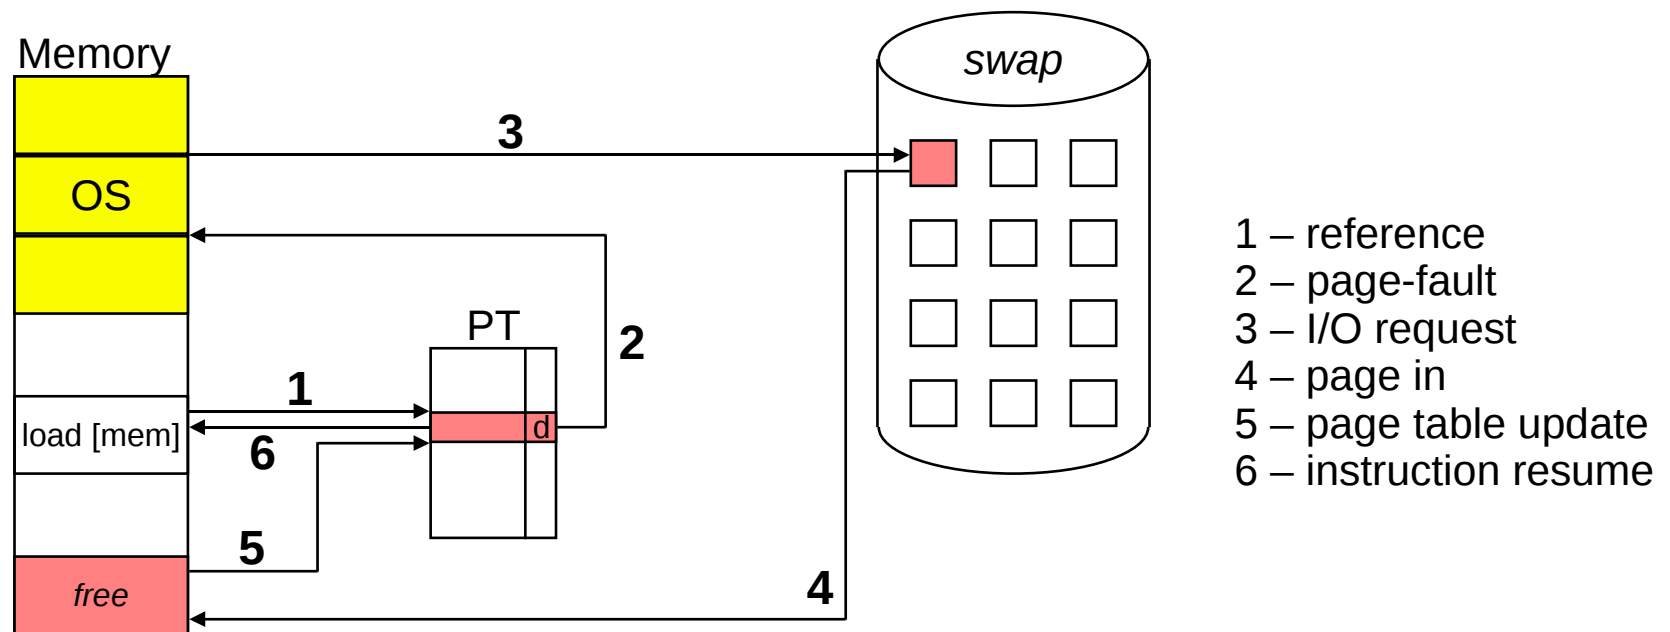
- Allows a process to be executed even if not completely loaded into memory
- Allows for processes to allocate more memory than the size of physical memory
- Can improve CPU utilization
- Can reduce swap overhead
- Can kill your system!

Swapping

- Processes can be temporarily suspended and the memory allocated to them is first copied to a secondary storage (i.e. disk) and then released to other processes (*swap out*)
 - Sleeping processes
 - Low-priority processes
- Such processes can be latter resumed by restoring their address spaces from the copy in the secondary storage (*swap in*)

Demand Paging

- Page-oriented swapping
 - Page table flag indicates whether the page is in memory or on disk
 - MMU triggers an exception (*page-fault*) whenever an absent page is accessed



Page-Fault Handling

Page-fault trap	1 μ s
Save context	10 μ s
Dispatch PF handler	1 μ s
Locate page on disk	50 μ s
Read page from disk	10ms
Waiting queue	0s
Seek	7ms
Latency	2ms
Transfer	1ms
Scheduler	15 μ s
Disk I/O completion	1 μ s
Save context	10 μ s
Dispatch disk I/O handler	1 μ s
Update page table	15 μ s
Ready queue waiting	0s
Scheduler	15 μ s
TOTAL	10,109ms

Demand Paging Performance

■ Formula

$$\text{eat} = (1 - p) \times \text{mat} + p \times \text{pft}$$

eat = effective access time

mat = memory access time

pft = page-fault handling time

p = page-fault probability

■ Example

mat = 50 ns

pft = 10 ms = 10.000.000 ns

eat = $(1 - p) \times 50 + p \times 10.000.000$

eat = $50 + 9.999.950 \times p$

p = 0.001 \rightarrow eat = 10 μ s


eat = 50 ns \rightarrow p < 0,000.005 (1 / 200.000)

Page Replacement


- Whenever necessary, OS selects pages to be moved out to disk and then make them available to processes
- Algorithm criteria
 - Minimize page-fault ratio
 - Minimize I/O
- Dirty bit
 - Each page status (modified or not) is kept in the corresponding entry in the page table
 - MMU automatically sets that bit whenever a page is modified (written)
 - Modified (dirty) pages must be written to disk before being reused

First-In First-Out (FIFO)

- Replace the page that has been longer in memory (e.g. pages are time-stamped)
- Implemented using a FIFO queue
- Example




- Increasing memory may not decrease *pf* rate



Optimal

- Replace the page that will not be used for the longest period of time
- Not implementable, for it relies on knowing the future
- Example



Least Recently Used (LRU)

- Uses the recent past as an approximation of the near future
- Replace the page that has not been used for the longest period of time
- Example

Accesses																12 pf			
7	0	1	2	0	3	0	4	2	3	0	3	2	1	2	0	1	7	0	1
7	7	7	2		2		4	4	4	0			1		1		1		
	0	0	0		0		0	0	3	3			3		0		0		
		1	1		3		3	2	2	2			2		2		7		

- Implementation
 - Time-stamp for each page
 - Linked-list of pages

LRU Approximations

- Reference bit
 - Each page is assigned a reference bit that is set by the MMU whenever the page is accessed
 - OS clears those bits periodically
 - Order of use is unknown
 - Target page is any with cleared reference bit
- Reference word
 - Additional reference bits that are shifted by the MMU
 - Target pages are those with smallest reference values
- Second chance
 - Pages are tracked by a circular FIFO list
 - If the pointed page has a clear reference bit, it is taken to be replaced
 - Otherwise, the bit is cleared and the pointer is adjusted to the next page

LRU Approximations

- Least Frequently Used (LFU)
 - Uses a reference counter for each page that is incremented by the MMU
 - Target page is the one with smallest counter value
 - Pages intensively accessed in the past, but no longer in use will take long to be replaced
- Reference and Modification bits
 - In addition to accesses, MMU marks pages that have been modified
 - Replacement order
 - Not-accessed, not-modified
 - Not-accessed, modified
 - Accessed, not-modified
 - Accessed, modified

Allocation of Frames

- Minimum set of frames
 - Instructions and operands may be scattered across several pages
 - Architecture-dependent
 - Instructions must be restarted after a page-fault
- Frames per process
 - Proportional to process size (i.e. memory footprint)
 - Equal to all processes
- Process interference
 - A process might cause the replacement of a page initially allocated to other process
 - A process may only replace its own pages

Thrashing

- A process is “thrashing” if it spends more time replacing pages than executing
- Causes
 - OS monitors CPU utilization and allows more processes in
 - If CPU utilization was low due to page-faults, increasing the number of processes in the system might cause thrashing
 - Thrashing only occurs if global page replacement (i.e. from other processes) is allowed
- Prevention
 - At any given time, a running process must have a set of pages available that fulfills its demands: its working set of pages (time x space locality)

Final Considerations

- Process load
 - On-demand
 - At-once to main memory
 - At-once to swap disk
- Page size
 - Large -> less page-fault, less I/O, less page tables
 - Small -> less internal fragmentation
- I/O results
 - Pages that will receive I/O results must be pinned-down
- Programming and code generation
 - Although virtual memory is functionally transparent to programs, memory access patterns might have big influence on it (e.g. matrices)